

REMARKS/ARGUMENTS

Claims 1-20 are pending. Claims 1-8, 12, 13, and 15-19 have been amended. The drawings have been corrected. No new matter has been introduced. Applicants believe the claims comply with 35 U.S.C. § 112.

Claims 1-3 and 11-14

Claims 1-3 and 11-14 stand rejected under 35 U.S.C. § 102(b) as being anticipated by Ofek et al. (US 5,889,935).

Applicants respectfully submit that independent claim 1 as amended is novel and patentable over Ofek et al. because, for instance, Ofek et al. does not teach or suggest a secondary mirror volume (S-VOL) group including multiple S-VOLs created as mirror images of said primary volume, and an S-VOL restoring unit configured to restore the data of a first S-VOL belonging to the S-VOL group with the data of a second S-VOL belonging to the S-VOL group depending on the type of an error that happens in the first S-VOL.

Ofek et al. discloses a mirroring system between storage systems, wherein a primary volume 22a (R1-VOL) and a secondary volume 50b (R2-VOL) form a pair, and data that is written into the R1-VOL is also written into the R2-VOL. If a failure occurs in the R1-VOL, the R1-VOL is restored using the data stored in the R2-VOL.

In the present invention as claimed, a plurality of secondary volumes (S-VOLs) are provided for one primary volume (P-VOL) and grouped together. If a failure occurs in an S-VOL, the failed S-VOL is restored using data stored in another S-VOL belonging to the same group as the failed S-VOL. In this way, it is possible to secure the availability of an S-VOL.

The Examiner states that the restoration between the R1-VOL and R2-VOL in Ofek et al. corresponds to the restoration between S-VOLs by an S-VOL restoring unit in the present invention. In Ofek et al., however, the R1-VOL and R2-VOL form a pair and there is no disclosure that a plurality of R2-VOLs grouped together are provided for the R1-VOL. Furthermore, Ofek et al. does not disclose that if a failure occurs in an R2-VOL, the failed R2-VOL is restored using data stored in another R2-VOL belonging to the same group as the

R2-VOL. In short, Ofek et al. does **not** teach or suggest an S-VOL restoring unit that restores data between S-VOLs grouped together, such that if a failure occurs in an S-VOL, the failed S-VOL is restored using another S-VOL belonging to the same group as the failed S-VOL.

For at least the foregoing reasons, claim 1, and claims 2-3 and 11 depending therefrom, are novel and patentable over Ofek et al.

Applicants respectfully submit that independent claim 12 as amended is novel and patentable over Ofek et al. because, for instance, Ofek et al. does not teach or suggest a secondary mirror volume (S-VOL) group including multiple S-VOLs created as mirror images of said primary volume, and restoring the data of a first S-VOL belonging to the S-VOL group with the data of a second S-VOL belonging to the S-VOL group depending on the type of an error that happens in the first S-VOL.

As discussed above, Ofek et al. discloses that the R1-VOL and R2-VOL form a pair and fails to teach that a plurality of R2-VOLs grouped together are provided for the R1-VOL. Furthermore, it is not disclosed that if a failure occurs in an R2-VOL, the failed R2-VOL is restored using data stored in another R2-VOL belonging to the same group as the R2-VOL. Thus, Ofek et al. does **not** teach or suggest restoring data between S-VOLs grouped together, such that if a failure occurs in an S-VOL, the failed S-VOL is restored using another S-VOL belonging to the same group as the failed S-VOL.

For at least the foregoing reasons, claim 12 and claims 13-14 depending therefrom are novel and patentable over Ofek et al.

Claims 4-10 and 15-20

Claims 4-10 and 15-20 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Ofek et al. in view of Kamiyama (US 5,893,139). The Examiner recognizes that Ofek et al. does not teach a monitor unit configured to monitor frequencies of accesses to a read-only volume and to restore the data of a volume where an error has happened by copying the data of an S-VOL with the lowest access frequency, and cites Kamiyama for allegedly providing the missing teachings.

Applicants respectfully assert that Kamiyama does not cure the deficiencies of Ofek et al., in that it also fails to teach or suggest a secondary mirror volume (S-VOL) group

including multiple S-VOLs created as mirror images of said primary volume, and an S-VOL restoring unit configured to restore the data of a first S-VOL belonging to the S-VOL group with the data of a second S-VOL belonging to the S-VOL group depending on the type of an error that happens in the first S-VOL, as recited in claim 1 from which claims 4-10 depend; and a secondary mirror volume (S-VOL) group including multiple S-VOLs created as mirror images of said primary volume, and restoring the data of a first S-VOL belonging to the S-VOL group with the data of a second S-VOL belonging to the S-VOL group depending on the type of an error that happens in the first S-VOL, as recited in claim 12 from which claims 15-20 depend.

Dependent claims 4-10 and 15-20 are patentable as being directed to additional features of the invention as well as by being dependent from allowable claims 1 and 12. Examples of those features include the frequency of access (claims 4, 5, and 7), the increments-volume (claims 6 and 7), and the spare volume (claim 8).

As to the frequency of access (claims 4, 5, and 7), Kamiyama discloses a method wherein while monitoring the frequency of access to data, the data to which the frequency of access is high is arranged in a storage having high access speed, and the data to which the frequency of access is low is arranged in a storage having low access speed. In this way, by rearranging data between storages on the basis of the frequency of access, the decrease in access performance of the entire system is prevented.

In contrast, in the claimed invention, while monitoring the frequencies of accesses to RO S-VOLs, if a failure occurs in an S-VOL, the S-VOL is restored using an RO S-VOL having the lowest frequency of access of the RO S-VOLs belonging to the same group as the failed S-VOL. In this way, by using an RO S-VOL having the lowest frequency of access, influence on processes using the RO S-VOLs can be reduced.

The Examiner asserts that Kamiyama discloses the use of an access management section for monitoring the frequency of access to data and rearranging in storages, and that it would have been obvious to combine Ofek et al. and Kamiyama to perform the restoration of a failed S-VOL using an RO S-VOL having the lowest frequency of access. In Kamiyama, however, while monitoring the frequency of access to data and not the frequency of access to an RO S-VOL, the data is rearranged between storages on the basis

of the frequency of access. Moreover, in Ofek et al. and Kamiyama there is no disclosure that the restoration of a failed S-VOL is performed on the basis of the frequencies of accesses to RO S-VOLs. Therefore, it would not have been obvious in view of Ofek et al. and Kamiyama to perform the restoration of a failed S-VOL using an RO S-VOL having the lowest frequency of access.

As to the increments-volume (claims 6 and 7), a cache 28 in FIG. 1 of Ofek et al. temporarily stores data with which to update the R1-VOL and R2-VOL, and if a failure occurs in the R1-VOL, the R1-VOL is restored using the R2-VOL as mentioned above. In contrast, in the claimed invention, updates to an RW S-VOL are stored in an increments-volume, and if a failure occurs in the RW S-VOL, the RW S-VOL is restored by updating an RO S-VOL belonging to the same group as the RW S-VOL with data stored in the increments-volume.

The Examiner asserts that the cache 28 of Ofek et al. corresponds to the increments-volume of this application and it would have been obvious to restore the RW S-VOL by updating an RO S-VOL with data stored in the increments-volume. Although certainly the cache 28 stores data with which to update the R1-VOL, in Ofek et al., it is not disclosed that if a failure occurs in the R1-VOL, the R1-VOL is restored *by updating the R2-VOL, rendered read-only, with data stored in the cache 28*. Therefore, it would not have been obvious in view of Ofek et al. and Kamiyama to restore the RW S-VOL by updating an RO S-VOL with data stored in the increments-volume.

As to the spare volume (claim 8), Ofek et al. disclose a method wherein if a failure occurs in the R1-VOL or the R2-VOL, restoration is performed using a spare. For example, if a failure is likely to occur in the R1-VOL, data from the R1-VOL is copied into a spare. After the copy finishes, the spare is used instead of the R1-VOL. Furthermore, for example, if a failure occurs in the R2-VOL, data from the R2-VOL forming a pair with the R1-VOL is copied into the spare. After the copy finishes, the spare is used instead of the R2-VOL. In either case, after the failed disk of the R1-VOL is replaced, the R1-VOL is restored by re-synchronizing the R1-VOL and the spare. In contrast, in the claimed invention, provided for an S-VOL is a spare S-VOL having read/write accesses thereto forbidden which belongs to the same group as the S-VOL and keeps the content at a certain time of the S-

VOL. If a failure occurs in the S-VOL, the failed S-VOL is restored using the spare S-VOL belonging to the same group as the failed S-VOL. In this way, a failed S-VOL is restored using not an RO S-VOL but a spare S-VOL having read/write accesses thereto forbidden, and thereby processes using the RO S-VOLs can be prevented from being affected.

The Examiner asserts that the spare of Ofek et al. corresponds to the spare S-VOL of this application and it would have been obvious to restore an S-VOL by using a spare S-VOL instead of an RO S-VOL. The spare of Ofek et al., however, does not store data beforehand like the spare S-VOL of this application, but data from the R1-VOL or the R2-VOL is copied into the spare when a failure occurs in the R1-VOL. After the copy finishes, the spare is used instead of the R1-VOL. That is, in Ofek et al., when a failure occurs, data from the R1-VOL or the R2-VOL needs to be copied into the spare. Thus, processes using the R1-VOL or the R2-VOL are affected. On the other hand, according to the configuration of the claimed invention, data from the S-VOL is stored beforehand in the spare S-VOL. Hence, data need not be copied into the spare S-VOL when a failure occurs. Thus, processes using the volumes can be prevented from being affected.

Moreover, in Ofek et al., after the failed disk of the R1-VOL is replaced, the R1-VOL is restored by copying data stored in the spare into the R1-VOL (re-synchronizing). However, in the re-synchronizing, the spare is not controlled to have read/write accesses thereto forbidden, but a HOST reads from/writes into the spare. Hence, the spare in this state does **not** correspond to *the spare S-VOL in the claimed invention that keeps the content at a certain time of an S-VOL and that is controlled to have read/write accesses thereto forbidden*. Therefore, it would not have been obvious in view of Ofek et al. and Kamiyama to restore an S-VOL by using a spare S-VOL instead of an RO S-VOL.

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance. The issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 650-326-2400.

Respectfully submitted,



Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
Attachments
RL:rl
60625377 v1

Amendments to the Drawings:

The attached sheet of drawings includes changes to Fig. 2. This sheet, which includes Fig. 2 replaces the original sheet including Fig. 2.

Attachment: Replacement Sheet
Annotated Sheet Showing Changes

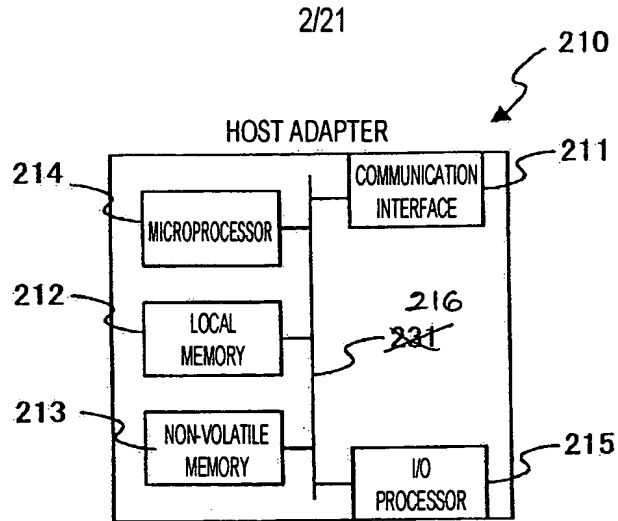


FIG. 2

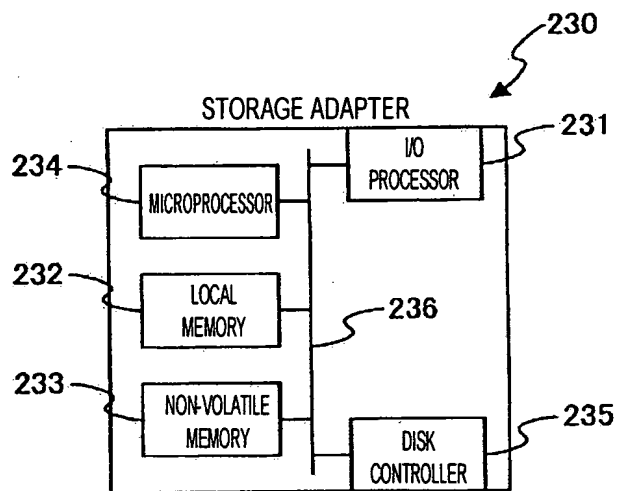


FIG. 3

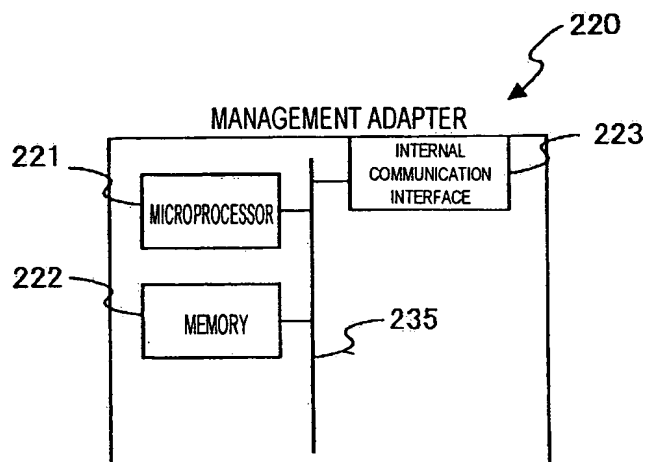


FIG. 4